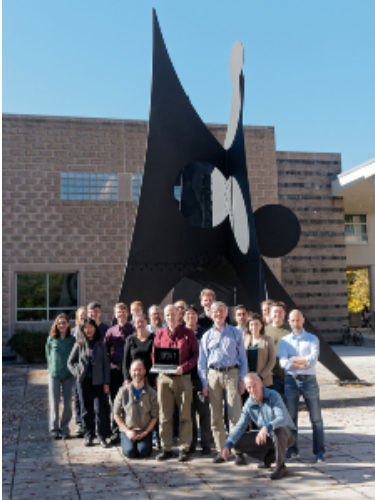


DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08



Logistics

Date

06 Nov 2018 – 08 Nov 2018

Expect the DMLT meeting to start in the morning of 06 Nov 2018 and conclude around lunchtime on 08 Nov 2018 to enable travel that afternoon.

That there will be a pre-DMLT SST meeting on 05 Nov 2018 .

Location

- Princeton

Teleconference

- Blue Jeans 925683803 (<https://bluejeans.com/925683803>, <http://ls.st/ski>)

Hotel / Per Diem

- There is a block booking at the [Nassau Inn](#). Group Code "Booking ID" is **24465**. Please contact [Yusra AlSayyad](#) for assistance.

Travel

- Princeton is most conveniently reached by flying to [Newark Liberty International Airport](#) (EWR) then taking an [NJ Transit](#) train to Princeton (change at Princeton Junction).
 - Note that the "Dinky" — the train from Princeton Junction to Princeton — will be suspended during the DMLT meeting. A substitute bus service will be provided by NJ Transit. Please allow extra time for your journey.
 - [More details from NJ Transit](#).
- It's also accessible from other New York area airports (La Guardia, JFK) or from Philadelphia International.
- The [Department of Astrophysical Sciences](#) has [travel advice](#).
- The DM-SST meeting will take place in the [Peyton Hall](#) "Dome Room".
- The DMLT meeting will take place in [Jadwin Hall](#), Joe Henry Room, Room 102.

Attendees

Traveling

- [John Swinbank](#)
- [Unknown User](#) (mbutler)
- [Unknown User](#) (xiuqin)

- [Leanne Guy](#)
- [Eric Bellm](#)
- [Unknown User \(gcomoretto\)](#)
- [Fritz Mueller](#)
- [Vaikunth Thukral](#)
- [Simon Krughoff](#)
- [Jeff Kantor](#) (cannot attend in person, will attend some sessions remotely from Tucson)
- [Wil O'Mullane](#)
- [Colin Slater](#)
- [Frossie Economou](#) (remote)
- [Tim Jenness](#)
- [Zeljko Ivezić](#)
- [Kian-Tat Lim](#)
- [Margaret Gelman](#)
- [Robert Gruendl](#)

Local

- [Yusra AlSayyad](#)
- [Robert Lupton](#)
- [Jim Bosch](#)


Agenda

Day 0: 2018-11-05			
DM-SST meeting. Agenda will be provided by Leanne Guy .			
Day 1: 2018-11-06			
Time	Topic	Chair	Notes
09:00	Welcome (and all my slides)	Wil O'Mullane	<ul style="list-style-type: none"> • Review outstanding actions • Review & adjust agenda

09:15	Impact of scope changes	Wil O'Mullane & Unknown User (xiuqin)	<ul style="list-style-type: none"> What portal development will happen, and when? How does this impact on other aspects of the DM system (e.g. who provides the UI for specifying and managing alert filters)? What other scope options are in the offing? Is further work to define new scope options required? <hr/> <p>Wil O'Mullane :</p> <ul style="list-style-type: none"> Slides in Wil O'Mullane 's master deck. We discussed whether some DM work can be pushed into the future as effectively a “no cost extension” to construction; the practicalities of this are not clear... <ul style="list-style-type: none"> In terms of what's acceptable to the agencies. And in terms of whether there will be staffing in the operations era to work on this. If there are people in your team/institute who would be interested in moving to AURA positions, please let Wil O'Mullane know. <ul style="list-style-type: none"> There are practical issues here, regarding which states are acceptable for people to reside, and whether there is office space available for them. And of course the downsides to breaking up groups. DM10 alert filters interface may actually represent a scope <i>increase</i> in O2C.03 in terms of providing a mini-broker interface. We believe that the current portal <i>is</i> sufficient for commissioning data releases (after finishing the closeout plan). We discussed increase staffing load resulting from making commissioning data widely available: this is expected to be significant. <p>Unknown User (xiuqin) :</p> <ul style="list-style-type: none"> LSST-based scale testing of Firefly has not been performed, but there is experience deploying it at scale in IPAC. A low-level of support is not enough to get fast turnaround on critical bug fixes. We discussed whether closeout work can focus on usability and performance enhancements rather than new feature development. <ul style="list-style-type: none"> There is work ongoing here, but the DMLT regards it as necessary to focus on LSST-specific work (which mainly concerns VO services). <p>Discussion:</p> <ul style="list-style-type: none"> There will be an LCR to cover the Portal descope. <ul style="list-style-type: none"> This will likely <i>also</i> include savings from the Data Facility and any other savings. And adjustments to the DM Science budget. And changes to the DM travel funds. Wil O'Mullane will be writing this, in discussion with Unknown User (xiuqin) . Will aim to get this to the CCB before the end of this year. There will be a presentation to the SAC on roughly the same timescale. This does not block the LCR. <p><input checked="" type="checkbox"/> All DMLT members: suggest changes to the Portal priority list if necessary. 09 Nov 2018</p>
10:00	CAOM & image metadata	Jim Bosch	<ul style="list-style-type: none"> Relationship between our image metadata model and CAOM. <hr/> <ul style="list-style-type: none"> No slides. Mapping the CAOM standard to existing DM technology & terminology. <ul style="list-style-type: none"> CAOM is becoming a widespread standard in the community; astronomical data centers are adopting it. It does not attempt to handle database records. Certainly appropriate for single-epoch data and “regular” coadds; extension to cell-based coadds is not obvious. Difficulties representing provenance in the CAOM model. <ul style="list-style-type: none"> e.g. no hierarchy. Do not expect to use CAOM as the “source of truth” for all DM provenance. We note that the CAOM system may not map obviously to LSST data, but that it has been used for other surveys. There is ongoing discussion between us and the IVOA folks. Conclusion: <ul style="list-style-type: none"> We will not further shoehorn the registry into looking like CAOM, but will maintain a separate structure which maps registry to CAOM. Export of CAOM from our data model is an explicit action (e.g. it requires running code).
10:30	Break (Refreshments Provided)		
11:00	The Science Data Model	Colin Slater	<ul style="list-style-type: none"> Provide an overview of work which has been undertaken in the DM-SST. DMLT to agree on direction of travel, deliverables, etc. <hr/> <ul style="list-style-type: none"> SDM Slides A single source of truth for both what columns “mean” in the scientific sense and how they are instantiated in a database. Proposal to move the contents of the cat package into a YAML file describing the “Science Data Model”. We discussed keeping the YAML definition of the SDM under change control. This is a minor point: it just needs somebody to define the procedure. We discussed whether YAML is an appropriate format for describing the database schema; the consensus is “yes”. Would ultimately like to use this YAML definition to describe the Butler registry. Would also like to generate the alert schema from the SDM. <ul style="list-style-type: none"> There may be some further work here to map to the hierarchical AVRO system. Next step is for this plan to be RFCed. <ul style="list-style-type: none"> At that point, it will become part of the DM baseline. And then management (ie, Wil O'Mullane) will need to figure out who is actually delivering the work.


12:00	"DPDDification"	Yusra AlSaiyyad	<ul style="list-style-type: none"> Review the tools which have been produced since the last meeting to turn pipeline outputs into something resembling the DPDD. Describe how this relates to work on the Data Model. <hr/> <ul style="list-style-type: none"> Notebook <code>writeObjectTable</code> (the precursor to transformation) currently runs in ~a minute per patch <ul style="list-style-type: none"> 45s to read in 15 <code>afwTables</code> from GPFS. <code>run()</code> takes 10s. 5s to write.. The DMLT agrees that Pandas/Parquet is an appropriate technology choice for the remainder of construction (at least). Inconclusive discussion around bitpacking of flags. <ul style="list-style-type: none"> Refer also to the "outstanding questions" in Yusra's notebook. Relationship to AP? <ul style="list-style-type: none"> May be necessary to go directly to database. But should also be possible to go to Parquet for debugging purposes. Requirement for round-tripping to the database. "Some careful thought about whether this can be brought to the AP side"; we should try hard to make this possible. Steps from here: <ul style="list-style-type: none"> Write up this proposal with a DRP-biased perspective. Ask the AP team to figure out which parts of this code are relevant (in conjunction with Yusra AlSaiyyad , Colin Slater , etc). We expect that the SDM will apply directly to AP. <p>✓ John Swinbank Arrange a meeting with Unknown User (cmorrison) , Yusra AlSaiyyad , Colin Slater , Eric Bellm to discuss SDM standardization in AP. 31 Jan 2019</p>
12:30	Lunch (Provided)		
13:30	Transition to Commissioning	Wil O'Mullane	<ul style="list-style-type: none"> "As discussed at the PST F2F meeting in September, DM should review the names of people going into commissioning" (requested by Leanne Guy) Not immediately clear whether this means permanent reassignment of staff, or those DM folks who are temporarily assigned to assist the Commissioning Team with specific activities. <hr/> <ul style="list-style-type: none"> Slides in Wil O'Mullane 's master deck. The commissioning support activities listed are <i>validation</i> of the DM system, but <i>verification</i> of the LSST system. None of the commissioning activities listed cover aspects of the system outside Science Pipelines. Concern expressed about the impacts of people being reassigned to commissioning on the rest of the DM team. <ul style="list-style-type: none"> Often the people who might be most effective with commissioning are also those most necessary for facilitating efforts within DM; concern expressed that this will have a disproportionate impact on the DM schedule. Detailed definition of the contents of the commissioning tests is with Leanne Guy and Keith Bechtol . Expectation that they will often involve repeating tests that have been performed within DM.
14:15	Transition to Operations	Wil O'Mullane	<ul style="list-style-type: none"> Early Ops funding is now available, and during FY19 (ie, this year) the ADs for both Data Facility and Science Operations (Margaret Gelman and Wil O'Mullane) are funded at 0.25 FTE. That's half an FTE coming out of DM management. How are we handling that? Similarly there's a total of ~15 FTEs funded across Data Facility and Science Ops during FY20. Please review: <ul style="list-style-type: none"> The plans and schedule for transitioning staff; The activities which the Operations Team will be carrying out with this effort, and how they relate to ongoing DM construction and the commissioning effort. <hr/> <ul style="list-style-type: none"> Slides in Wil O'Mullane 's master deck. FY19 includes 25% of Wil O'Mullane & Margaret Gelman , as well as Phil Marshall on science performance; there are no milestones in this. Some milestones (e.g. ops rehearsals) are effectively duplicated between DM and pre-operations funding; where possible, they will migrate from DM to pre-ops. There's a worry that commissioning data <i>may</i> become available to the public more quickly than we currently plan; DM should be ready to scale up to address this. <ul style="list-style-type: none"> And a feeling that simply making the data available for download will not adequately address this need. But we acknowledge the cost and schedule impacts of this. There's also a concern about what level of end-user support is implied by this. How do we handle folks being required in DM <i>and</i> commissioning <i>and</i> pre-ops? <ul style="list-style-type: none"> This is a matter of ongoing planning. There may be some overlap/double counting. Discussion of how open the access to commissioning data & facilities should be, balancing getting input on commissioning from members of the community with the support load and "chaos" of wide access. <p>✓ Wil O'Mullane — coordinate the writing of a memo describing what community DM can support during commissioning. 03 Dec 2018</p>
15:00	Break (Refreshments Provided)		


15:30	Management of externally contributed packages	Wil O'Mullane	<ul style="list-style-type: none"> The proposal is to develop a policy for handling packages which have been developed externally and which their authors offer up for inclusion in pipeline processing, the Science Platform environment, or elsewhere in the DM system. Obvious examples might be scientific algorithms contributed by the community. There is history of external users refusing to make contributions like this due to the demands of DM engineering (code quality, review, tests, etc...). <hr/> <ul style="list-style-type: none"> Slides in Wil O'Mullane 's master deck. We did not converge on a requirement for simply installing software for pure end user convenience: they can be e.g. pip installed. But we do identify this as a potential way forward for software from science collaborations which might migrate into the LSST codebase. <ul style="list-style-type: none"> But there's no urgency to address this now. Discussion of Erin Sheldon/ngmix specifically: <ul style="list-style-type: none"> Want to run this in Data Release Production, but not maintain it ourselves. How do we incorporate this into our pipeline without forcing them and/or us to stop development or rewrite it. This requires definition of an interface. <p>☑ Wil O'Mullane with Kian-Tat Lim , Frossie Economou , Unknown User (gcomoretto) — draft a policy for including external code into data release production processing. 03 Dec 2018 see</p> <div style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  DM-17964 - Jira project doesn't exist or you don't have permission to view it. </div>
16:00	Product Tree and Document Tree	Unknown User (gcomoretto)	<ul style="list-style-type: none"> Based on the modeling work done before the review this year, review proposed product tree, components characterization and relation with the document tree <hr/> <ul style="list-style-type: none"> "Inside pipelines, there are more or less five products" — not every Git repository or software package is a product. "SW Products can depend on other SW products without containing them" — so there is a SW product that contains e.g. the Butler, which can be depended upon by other products. SW Products are the unit of release. Dependency relationships happen between SW products, rather than between SW packages. <p>☑ All DMLT: Review product tree in LDM-294 and provide feedback/corrections to Unknown User (gcomoretto) . 03 Dec 2018</p> <p>☑ Unknown User (gcomoretto) : to produce a technical note including all products from the product tree and their characterization. 07 Jan 2019</p>
16:30	Test Approach Using Jira	Unknown User (gcomoretto)	<ul style="list-style-type: none"> How to create test specifications and perform test runs with Adaptavist Test Management in Jira. <hr/>
17:00	Close		
Day 2: 2018-11-07			

09:00	Middlewar e & Workflow	Margaret Gelman & Fritz Mueller	<ul style="list-style-type: none"> Where are we on the selection & deployment of the workflow management system? What's the development timeline for next generation middleware (Butler Gen 3, PipelineTask)? What's the overall roadmap for middleware and workflow over the next ~4 years? <hr/> <ul style="list-style-type: none"> Slides: <ul style="list-style-type: none"> Fritz MuellerMiddleware Update DMLT Nov 2018.pptx Margaret GelmanWorkflow Futures.pdf The Data Facility can execute based on whatever version of the middleware the developers are using. <ul style="list-style-type: none"> There's no requirement for pipeline developers to support Gen2 from the LDF point of view. Most Data Facility work going forward is based on Pegasus; there is minimal ongoing support for DESDM. Assuming availability of pipeline code, the LDF predicts that they could run pipelines in a "sustained processing" mode based on Butler G3 and Pegasus in mid-2019. Two possible goals for BG3 priority: <ul style="list-style-type: none"> Support for obs_lsst (to make RHL & Merlin's life easier) Or to convert code to PipelineTask to enable execution at scale of e.g. HSC on the Data Facility. The consensus is that the latter is the priority; agreed to prioritise the conversion of ci_hsc Tasks to PipelineTasks until end of Jan 2019 per Fritz Mueller's recommendation. <ul style="list-style-type: none"> This would also meet Frossie Economou's immediate validate_drp use case. We agreed that <i>temporarily</i> abandoning the shared-nothing model for execution might enable faster development. We should <i>also</i> prepare for "plan B" by assembling a "mini-working-group" to consider wholesale technological change (mini-WG to consist of at least Kian-Tat Lim & Simon Krughoff; not to involve folks who are busy with the ci_hsc conversion). <div style="border: 1px solid #fde725; padding: 10px; margin-top: 10px;"> <p> Further middleware discussion</p> <p>A further, unscheduled, discussion of middleware took place at 09:00 on the following morning. Notes from it are recorded below. They supersede the above, and render some action items previously recorded here obsolete.</p> </div> <div style="margin-top: 10px;"> <input checked="" type="checkbox"/> John Swinbank — designate and/or start a recruitment process for a "systems programmer" to act as long term middleware owner. 01 Mar 2019 </div>
10:20	Long-term release support	Leanne Guy	<ul style="list-style-type: none"> Requirements for back-porting bug fixes to release branches, in support of science collaborations, commissioning, etc. <hr/> <ul style="list-style-type: none"> Slides <div style="margin-top: 10px;"> <input checked="" type="checkbox"/> Kian-Tat Lim — finalise and document deprecation procedure (ie, implement RFC-213, with whatever modernization or updates are necessary). 03 Dec 2018 </div>
10:40	Break (Refreshments Provided)		
11:00	Release Process	Unknown User (gcomoretto)	<ul style="list-style-type: none"> Proposed changes to the release process. <hr/> <ul style="list-style-type: none"> Assert that a distribution that consists only of "the binaries" necessary for release to operations is required. <ul style="list-style-type: none"> It's not clear that a consensus in the room agreed with this point. <div style="margin-top: 10px;"> <input type="checkbox"/> </div> <div style="margin-top: 10px;"> <input checked="" type="checkbox"/> Leanne Guy & Unknown User (gcomoretto) — propose a procedure for release maintenance, including back-porting of bug fixes to stable releases. 19 Apr 2019 </div>

11:30	Review of Calibration plans	Robert Lupton	<ul style="list-style-type: none"> What raw data is being taken? What products are being generated? What are the possible CPP execution periodicities <hr/> <ul style="list-style-type: none"> Atmospheric absorption data structure TBD; may be a lookup table, for example. Need to capture generation of a distortion model; this should be part of calibration products (ie, capture the output of Jointcal for use in pipelines processing). Outstanding question is which calibration data has to be fed through the prompt processing system for on-the-fly adjustments of the observatory configuration. <ul style="list-style-type: none"> The answer is "no" — where necessary, they can be deployed as part of T&S software on the mountain, not in a DM execution framework. We expect the Commissioning Cluster will not be reliably available during operations, so would not be a good home for this. Worries were expressed that the future capabilities, and requirements for capabilities, provided by a Commissioning Cluster-like service are unclear. <p>✓ John Swinbank / Robert Lupton — Ensure that the plans for how calibration products pipelines will be executed during operations are clear, including e.g. executing Jointcal to produce a distortion model for ingestion into the science pipelines. 08 Jan 2019</p> <p>✓ Wil O'Mullane — Add a discussion of future requirements for Commissioning Cluster-like capabilities during the operational era to the AuxTel workshop in January. 26 Nov 2018</p> <p>✓ Robert Lupton — Upload slides on calibration plans to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018</p>
12:30	Lunch (Provided)		
13:30	Summit and Base Data Center Status and Planning	Jeff Kantor	<ul style="list-style-type: none"> Brief overview of Networks, Summit and BDC construction status and move-in schedule Identification of planned deployments of DM equipment to BDC, visits, tests/rehearsals in FY19 (and IT support required) <hr/>
14:30	DBB and Consolidated DB SLA	Kian-Tat Lim	<ul style="list-style-type: none"> Are there ways of providing greater uptime and lower latency using rolling upgrades, schema evolution with backfill, planned Observatory maintenance windows and daytime maintenance, etc.? What are the features and timeline for DBB? <hr/> <ul style="list-style-type: none"> Discussion about how databases are represented within the DBB. <ul style="list-style-type: none"> The DBB has to know about different types of databases; an Oracle DB is replicated differently from other technologies, for example. Qserv is not part of the DBB. <ul style="list-style-type: none"> Even Qserv replication is entirely separate from the DBB. How are spherical geometry queries performed within the consolidated DB? This remains unresolved. <ul style="list-style-type: none"> ADQL support is required for performing TAP queries. Multiple possibilities exist: externally-computed pixelization indexes, Google S2, Oracle Spatial product. Should incorporate availability constraints into test plans for the DBB. The TAP service will not be required to interface to the "live" PPDB, just the equivalent information in the Consolidated Database that derives from the PPDB. This implies that this data gets to the Consolidated DB with whatever latency would meet the requirements for making that data available through the TAP service. The baseline for doing so at this point is ≤ 24 hours but Eric Bellm is working on usecases that could result in tightening that. (On the other hand, it was mentioned that DoD concerns may motivate keeping this at ≥ 24 hours.) <p>✓ Tim Jenness — provide astro_metadata_translator for obs_Isst by December, but at lower priority than PipelineTask conversion of ci_hsc. 03 Dec 2018</p> <p>✓ Kian-Tat Lim — draft an SLA document for each enclave and services within it. 07 Jan 2019 Now</p> <div style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  DM-17836 - Jira project doesn't exist or you don't have permission to view it. </div> <p>✓ Wil O'Mullane & Zeljko Ivezic — clarify whether any changes to the baseline for L1PublicT are required. 03 Dec 2018</p>
15:00	Break (Refreshments Provided)		
15:30	Review of L2 Milestones during calendar 2019	John Swinbank	<ul style="list-style-type: none"> Which milestones are scheduled? Who is responsible for them? What other milestones might be added? <hr/>

- For each of the below, we need to confirm:
 - The due date
 - The responsible person
 - The contents of the milestone
 - What support is needed from other teams to meet this milestone; cross-check with LDM-564
- Actions on milestone owners to:
 - Update LDM-503 to show a brief description of each milestone.
 - Ensure that prerequisites described in LDM-564 are correct.

Milestone	Name	Date	Responsible	Relevant tickets	Notes
LDM-503-06	ComCam Interface Verification	2018-11-30	Margaret Gelman	 DM-16074 - Jira project doesn't exist or you don't have permission to view it.	Pending availability of DAQ at NCSA L1 test stand.
LDM-503-07	Camera Data Processing	2018-11-30	John Swinbank	 DM-16072 - Jira project doesn't exist or you don't have permission to view it.	
LDM-503-09	Ops Rehearsal #1	2019-04-30 (per LCR-1488)	Robert Gruendl	 DM-16075 - Jira project doesn't exist or you don't have permission to view it.	

LDM-503-09a	Pipelines Release Fall 2018	2018-11-30	John Swinbank	 DM-16076 - Jira project doesn't exist or you don't have permission to view it.	
LDM-503-08	Spectrograph Data Acquisition	2019-01-04	Unknown User (mbutler)	 DM-16455 - Jira project doesn't exist or you don't have permission to view it.	Should block on pipeline from Merlin. Output products to the Data Backbone.
DLP-552	LSST Software Release 9.1 Complete, Ready for Full Camera	2019-09-03	Wil O'Mullane	None.	We assert this means a Science Pipelines release, not other products. Tie it to LDM-503-11b.
LDM-503-10	DAQ Validation	2019-09-25	Unknown User (mbutler)	 DM-16193 - Jira project doesn't exist or you don't have permission to view it.	Prereq. for run from ComCam DAQ on mountain to NCSA. ComCam hardware availability will be tied to this test date.
LDM-503-11	Ops Rehearsal #2	2019-09-25	Robert Gruendl	 DM-16195 - Jira project doesn't exist or you don't have permission to view it.	Fleshed out further after rehearsal #1.

LDM-503-11b	Pipelines Release Fall 2019	2019-09-25	John Swinbank	 DM-16197 - Jira project doesn't exist or you don't have permission to view it.	
LDM-503-10b	Large Scale CCOB Data Access	2019-09-25	Unknown User (mbutler)	 DM-16194 - Jira project doesn't exist or you don't have permission to view it.	Pulling data from test stand at SLAC.
LDM-503-11a	ComCam Ops Readiness	2019-09-30	Robert Gruendl	 DM-16196 - Jira project doesn't exist or you don't have permission to view it.	This is an aggregation of previous tests, demonstrating that sufficient to operate ComCam have passed.
LDM-503-12	Ops Rehearsal #3	2019-11-26	Robert Gruendl	 DM-16198 - Jira project doesn't exist or you don't have permission to view it.	Fleshed out further after rehearsal #2.
DM-STAFF	Staffing Checkpoint	2019-12-02	Wil O'Mullane	None.	This will be satisfied by issuing a document; it does not need to go through the regular DM testing procedure.

			<div><div><div><div><div><div></div><div>John Swinbank — add DM-external prereqs for L2 milestones to LDM-503. 17 Dec 2018</div></div></div><div><div><div></div><div>John Swinbank , working with milestone owners — ensure appropriate prerequisites for all DM L2 milestones are captured in PMCS. 17 Dec 2018</div></div></div><div><div><div></div><div>Wil O'Mullane together with Gregory Dubois-Felsmann — define a 2019 era L2 milestone for demonstrating making data available in the LSP. 17 Dec 2018</div></div></div><div><div><div></div><div>Wil O'Mullane — add a milestone for the availability of a DAQ at the NCSA L1 test stand. 17 Dec 2018</div></div></div><div><div><div></div><div>Margaret Gelman — Prepare for LDM-503-06 by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 29 Mar 2019</div></div></div><div><div><div></div><div>John Swinbank — Prepare for LDM-503-07 by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 30 Nov 2018</div></div></div><div><div><div></div><div>Robert Gruendl — Prepare for LDM-503-09 i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div><div><div><div></div><div>John Swinbank — Prepare for LDM-503-09a by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 30 Nov 2018</div></div></div><div><div><div></div><div>Unknown User (mbutler) — Prepare for LDM-503-08 by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 17 Dec 2018</div></div></div><div><div><div></div><div>Unknown User (mbutler) — Prepare for LDM-503-10 by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div><div><div><div></div><div>Robert Gruendl — Prepare for LDM-503-11 by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div><div><div><div></div><div>John Swinbank — Prepare for LDM-503-11b by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div><div><div><div></div><div>Unknown User (mbutler) — Prepare for LDM-503-10b by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div><div><div><div></div><div>Robert Gruendl — Prepare for LDM-503-11a by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div><div><div><div></div><div>Robert Gruendl — Prepare for LDM-503-12 by i) ensuring that the description and comments in LDM-503 provide a concise summary of the aims and methodology of the milestone, and ii) ensuring that LDM-564 provides a complete list of prerequisites for this milestone. 25 Feb 2019</div></div></div></div></div></div>
16:00	Plans for Operations Rehearsal	Robert Gruendl	<div><div><div>• Update on the plans for, and timing of, the early-2019 Operations Rehearsal (LDM-503-09).</div></div><div><div><div>• Agreed to redefine the Ops Rehearsal such that it can move ahead without being tied it to the slipping Auxtel schedule.</div></div></div><div><div><div>Robert Gruendl — update LDM-643 describing plans for Ops Rehearsal #1 and circulate. 03 Dec 2018</div></div></div></div>
17:00	Close		
Evening	Group dinner at 287 Hamilton Ave. Please provide your \$34 dinner per diem to Wil O'Mullane .		
Day 3: 2018-11-08			

09:00	Planning for Gen 3 Middleware	Fritz Mueller , Jim Bosch et al.	<ul style="list-style-type: none"> • Unscheduled discussion, based partly on notes that Jim Bosch provided on Slack. • Arrived at the consensus that the plan to deliver a complete BG3 and PipelineTask conversion of all tasks in ci_hsc by end January 2019, arrived at yesterday, is unrealistic, given the other time commitments of major players. • Agreed on a variant of "option D" in Jim Bosch's notes, above. Specifically: <ul style="list-style-type: none"> ◦ BG3 will be integrated with the existing CmdLineTask framework. ◦ BG2 will be retired without waiting for all CmdLineTasks to be retired. ◦ CmdLineTasks may never be fully retired. As new tasks are written, or old tasks are refactored to produce pipelines which are more "LDM-151-like", they will be implemented as PipelineTasks, but no drive will be scheduled to complete the conversion of other CmdLineTasks. ◦ The Data Facility confirms that they will be able to support an execution environment capable of driving both PipelineTask and CmdLineTask indefinitely. • Success for the current work will be declared when BG3 has been fully adopted. At this point: <ul style="list-style-type: none"> ◦ Conversion of remaining tasks to the PipelineTask system (on an as-needed basis) becomes fully the responsibility of the Science Pipelines groups. ◦ A long-term maintainer for the BG3 codebase must be found, either as a new hire or (potentially) from within the ranks of the Data Facility. • Simon Krughoff volunteered effort to assist in converting code of particular relevance to SQuaRE to PipelineTasks, likely starting with ProcessCcdTask. There was some follow-up discussion about having his expertise best deployed on integration with obs packages instead; this should be included in Fritz Mueller's plan (below). <p>✓ Fritz Mueller — present a timeline for integration of the CmdLineTask framework with Butler Generation 3, and the subsequent deprecation of the BG2 system. 28 Jan 2019</p>
10:10	Level 3 Milestones	John Swinbank	<ul style="list-style-type: none"> • We're lacking L3 milestones in PMCS which describe the availabilities of services and capabilities which we know are coming, particularly later in construction. • Obvious examples include: <ul style="list-style-type: none"> ◦ Butler Gen 3 / PipelineTask as the regular production environment; ◦ Roll-out of the Pegasus WMS (or some other WMS); ◦ Data back bone capabilities; ◦ There are <i>no</i> DAX milestones beyond the end of November 2018; ◦ There are five SQuaRE milestones total, all of which refer to notebooks (and which are not necessarily SQuaRE deliverables — see below). • In addition, some existing milestones seem unclear about who is delivering what. For example: <ul style="list-style-type: none"> ◦ "DM-SUIT-16: Commissioning DAC" — is that really a SUIT deliverable? ◦ "DM-SQRE-5: Notebook service ready for general science" — is that really a SQuaRE deliverable? • I assert: <ul style="list-style-type: none"> ◦ We should have milestones describing the delivery of effectively everything described in LDM-148. <ul style="list-style-type: none"> ▪ Are there other things which are <i>not</i> in LDM-148 which we need to track? ◦ "Standing up a service" milestones are all ? for the Data Facility, and should be dependent on prerequisite milestones for software delivery. ◦ We should have brief (~few sentence) description for every milestone. <ul style="list-style-type: none"> ▪ We might consider having a more formal verification procedure, which would relate successful completion of these milestones to test execution. ▪ I suspect that such a procedure would lose us more in time & overheads than it gains us, but I am open to being convinced otherwise. • Can we produce a revised set of milestones? When is an appropriate due date — January 2019? <hr/> <ul style="list-style-type: none"> • Agreed to target the due date of this work as February 2019. • Agreed further that it should be driven by the product tree (not just LDM-148). Note that: <ul style="list-style-type: none"> ◦ This means work can't usefully begin until the product tree has been finalised (see discussion 2018-11-06 at 16:00). <ul style="list-style-type: none"> ▪ We note that the product tree is never really "final", as it will continue to evolve throughout construction, but we expect a substantially revised and updated version by the end of this year which will form the basis for this work. ◦ Some milestones are <i>not</i> directly related to DM products. These might include delivery of documentation or supporting services. For this reason (and for reasons of visibility/transparency, and because the PMCS will remain the source of truth regarding milestones) we will not associate milestones with products in MagicDraw. • Expect the minimal set of milestones will be availability of code to run a service (i.e. a software artefact) and the availability of the service itself (i.e. that software deployed at the Data Facility). In some cases (e.g. Science Pipelines) it will make sense to have multiple intermediate milestones for software delivery. <p>✓ John Swinbank — Following product tree updates, circulate to all T/CAMs a spreadsheet for collecting L3 milestones. 17 Dec 2018</p> <p>Actions on individual T/CAMs to update milestones in their WBS are captured at DM Technical Managers (T/CAMs).</p>
10:25	Break (Refreshments Provided)		

10:55	Plans for S19	John Swinbank	<ul style="list-style-type: none"> T/CAM or lead for each team spent 10 minutes outlining development priorities for S19. <ul style="list-style-type: none"> DM Science (Leanne Guy) Architecture (Kian-Tat Lim) Alert Production (John Swinbank) Data Release Production (Yusra AlSayyad) Science User Interface (Unknown User (xiuqin)) Data Access and Database (Fritz Mueller) LDF (Margaret Gelman) Base & Networks (Jeff Kantor) SQuaRE (Frossie Economou) Feedback from DMLT members on the plans heard earlier. Finalize cross-team priorities. Agree development plans for S19.
			<ul style="list-style-type: none"> DM Science: <ul style="list-style-type: none"> Documents for the LSP review are listed in the charge, and will be delivered to reviewers two weeks in advance of the view. Reviewers will cover both science and technical themes. Alerts Key Numbers study is not re-defining or deriving key numbers, but rather making them available to the community with adequate context. Not yet clear where 200 GB DESC Qserv test dataset will actually be hosted. Architecture: <ul style="list-style-type: none"> Sizing model work likely to happen as part of the LDF-operations funding; Arch ready to act in an advisory role. No due date. Data Access and Database: <ul style="list-style-type: none"> It is a DAX (Colin Slater) deliverable to ensure there's code for demonstrating that pipelines output matches the Science Data Model (but not for resolving discrepancies). Data Facility: <ul style="list-style-type: none"> Ongoing DESC DC2 processing based on informal discussions at LSST2018. This was news to most of the DMLT, but we agreed that it was a positive step. Wil O'Mullane is the lead organizer for the AAS demo session; he will coordinate necessary resources with LDF etc; Robert Lupton can send suggestions for scale testing / widespread access to him(!). Base & Networks: <ul style="list-style-type: none"> Jeff Kantor was not available. <p> <input checked="" type="checkbox"/> Leanne Guy — Upload slides on S19 planning to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018 </p> <p> <input checked="" type="checkbox"/> John Swinbank — Upload slides on S19 planning to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018 </p> <p> <input checked="" type="checkbox"/> Yusra AlSayyad — Upload slides on S19 planning to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018 </p> <p> <input checked="" type="checkbox"/> Unknown User (xiuqin) — Upload slides on S19 planning to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018 The information is in the slides "SUITdescopelmpact.pptx". You can also see here. Portal priority list </p> <p> <input checked="" type="checkbox"/> Fritz Mueller — Upload slides on S19 planning to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018 </p> <p> <input checked="" type="checkbox"/> Frossie Economou — Upload slides on S19 planning to DM Leadership Team Face-to-Face Meeting, 2018-11-06 to 08. 19 Nov 2018 </p>
12:30	Review action items & plans for next meeting	Wil O'Mullane	<ul style="list-style-type: none"> Next meeting 26–28 February in Tucson. Then May 21–23 May at NCSA. Then November 5–7 at SLAC. Next DMLT telecon 2018-11-26. <p> <input checked="" type="checkbox"/> Wil O'Mullane : Confirm dates of the November 2019 DMLT with the whole group. </p>
13:00	Close		

Pre-Meeting Planning

Suggested topics for discussion

Topic	Requested by	Time required (estimate)	Notes	Included in schedule above
Review of the DM people transitioning to commissioning	Leanne Guy		As discussed at the PST F2F meeting in September, DM should review the names of people going into commissioning	
Milestone test plans	Kian-Tat Lim		Continue work on content and dates for upcoming milestones	
Descope options	Kian-Tat Lim			

DBB and Consolidated DB SLA	Kian-Tat Lim		Are there ways of providing greater uptime and lower latency using rolling upgrades, schema evolution with backfill, planned Observatory maintenance windows and daytime maintenance, etc.?	
Review of Calibration plans	Kian-Tat Lim		What raw data is being taken? What products are being generated? What are the possible CPP execution periodicities?	
Unusual Commissioning modes	Kian-Tat Lim		What things that we normally think of as slowly-changing (e.g. camera geometry, voltage setpoints) might not be during Commissioning?	
Review of L2 milestones due during calendar 2019	John Swinbank	1 hour		
Review of plans for S19 development cycle	John Swinbank	2 hours		
Review impact of portal scope changes	John Swinbank / Wil O'Mullane	1.5 hours	<ul style="list-style-type: none"> What portal development will happen, and when? How does this impact on other aspects of the DM system (e.g. who provides the UI for specifying and managing alert filters)? 	
Status of DPDDification	John Swinbank	0.5 hours	<ul style="list-style-type: none"> Update on work started at the May DMLT F2F. 	
LSST Data Model	John Swinbank	1 hour	<ul style="list-style-type: none"> Work has been ongoing in the DM-SST to develop a formal DM data model and to refine the way in which LSE-163 (DPDD), the cat package, LDM-153 (the baseline schema) are generated and managed. This should be presented to and agreed by the DMLT. 	
Transition to operations	John Swinbank	1 hour	<ul style="list-style-type: none"> Early Ops funding is now available, and during FY19 (ie, this year) the ADs for both Data Facility and Science Operations (Margaret Gelman and Wil O'Mullane) are funded at 0.25 FTE. That's half an FTE coming out of DM management. How are we handling that? Similarly there's a total of ~15 FTEs funded across Data Facility and Science Ops during FY20. What are the plans and schedule for transitioning staff? 	
Product Tree and Document Tree	Unknown User (gcomoretto)	0.5 hour	<ul style="list-style-type: none"> Based on the modeling work done before the review this year, review proposed product tree, components characterization and relation with the document tree 	
Test Approach Using Jira	Unknown User (gcomoretto)	0.5 hour	<ul style="list-style-type: none"> How to do test specs and test runs with ATM in Jira 	
Summit and Base Data Center Status and Planning	Jeff Kantor	1 hour	<ul style="list-style-type: none"> Brief overview of Summit and BDC construction status and move-in schedule Identification of planned deployments of DM equipment to BDC, visits, tests/rehearsals in FY19 (and IT support required) 	
Workflow management	Leanne Guy	1 hour	<ul style="list-style-type: none"> Where are we on the decision and implementation of a WMS? 	
L3 milestones	John Swinbank / Wil O'Mullane	0.5 hour	<ul style="list-style-type: none"> We're lacking L3 milestones in PMCS which describe the availabilities of services and capabilities which we know are coming, particularly later in construction. Obvious examples include: <ul style="list-style-type: none"> Butler Gen 3 / PipelineTask as the regular production environment; Roll-out of the Pegasus WMS (or some other WMS); Data back bone capabilities. In addition, some existing milestones seem unclear about who is delivering what. For example: <ul style="list-style-type: none"> "DM-SUIT-16: Commissioning DAC" — is that really a SUIT deliverable? "DM-SQRE-5: Notebook service ready for general science" — is that really a SQuaRE deliverable? We should have milestones describing the delivery of effectively everything in LDM-148. The half hour time estimate is for Wil O'Mullane to state the above, set expectations, and develop a timeline for delivery. If we actually start creating milestones during the meeting, it could take arbitrarily long. 	
Management of externally contributed packages	Wil O'Mullane	0.5 hour	<ul style="list-style-type: none"> The proposal is to develop a policy for handling packages which have been developed externally and which their authors offer up for inclusion in pipeline processing, the Science Platform environment, or elsewhere in the DM system. Obvious examples might be scientific algorithms contributed by the community. There is history of external users refusing to make contributions like this due to the demands of DM engineering (code quality, review, tests, etc...) 	
Next generation middleware timeline	Fritz Mueller	0.5 hour	<ul style="list-style-type: none"> When do we expect BG3, PipelineTask, etc to be actively being used by developers? What's the roadmap for getting there? 	
Early considerations on Release process	Unknown User (gcomoretto)	0.5 hour	<ul style="list-style-type: none"> The aim is to share some preliminary considerations on the release process, taking in account the actual approach and giving a look on how this can evolve. 	(albeit in a potentially tightly squeezed slot...)

• Attached Documents

File	Modified
PDF File Kantor Networks and Base.pdf	Nov 06, 2018 by Jeff Kantor
PDF File Product Tree and Document Tree DMLT-F2F Nov 2018.pdf	Nov 06, 2018 by gcomoretto
PDF File Test Approach DMLT-F2F Nov 2018.pdf	Nov 06, 2018 by gcomoretto
PDF File long term release support -leanneguy.pdf	Nov 06, 2018 by Leanne Guy
PDF File OPS_Rehearsals_DMLT_20181107.pdf	Nov 07, 2018 by Robert Gruendl
PDF File Release Process DMLT-F2F Nov 2018.pdf	Nov 07, 2018 by gcomoretto
JPEG File DMLTNov2018.jpg	Nov 07, 2018 by Wil O'Mullane
PDF File DBB+CDB KTL DMLT 2018-11-07.pdf	Nov 07, 2018 by Kian-Tat Lim
Microsoft Powerpoint Presentation Middleware Update DMLT Nov 2018.pptx	Nov 07, 2018 by Fritz Mueller
PDF File Arch S19 Plans.pdf	Nov 08, 2018 by Kian-Tat Lim
Microsoft Powerpoint Presentation SUITdescopelImpact.pptx	Nov 08, 2018 by xiuqin
PDF File Workflow Futures.pdf	Nov 08, 2018 by Margaret Gelman
PDF File LDF Plan S19.pdf	Nov 08, 2018 by Margaret Gelman
PDF File 2018-11-08 — AP S19 Plans.pdf	Nov 13, 2018 by John Swinbank
PDF File lupton.pdf	Nov 13, 2018 by Robert Lupton
PDF File DRP-S19-Activities-Actual.pdf	Nov 13, 2018 by Yusra AlSayyad
PDF File dmlt_sdm.pdf	Nov 14, 2018 by Colin Slater
PDF File DAX S19 Plans.pdf	Nov 26, 2018 by Fritz Mueller
PDF File DM-SST Plans S19.pdf DM-SST Plans S19	Nov 26, 2018 by Leanne Guy
PDF File dmlt_2018_11.pdf	Dec 10, 2018 by Frossie Economou

[Download All](#)